DOE Center of Excellence Performance Portability Meeting Agenda (Draft - effective 3/23/2016)

Session	Day One	Speaker/Topic	Affiliation	Title
	7:30	Registration/Coffee/Light Breakfast		Mingle
Session chair:		**Overviews**	**Each Center of Excellence to give an overview. Projects, how help is supported, how vendors are integrated, etc…**	
Rob Neely	8:15	Welcome/Kickoff		
	8:30	Straatsma	ORNL	Summit COE / CAAR Overview
	8:40	Deslippe	LBL	NERSC-8 COE / NESAP Overview
	8:50	Neely	LLNL	Sierra COE Overview
	9:00	Nam/Glass/Dawson	LANL	Trinity COE Multi-lab Overview
	9:15	Riley	ANL	ANL COE Overview
	9:25	Riley	ANL	HPCOR Workshop Recap
	9:35	Still	Multi-lab	ECP Application Overview and Criteria
	9:50	BREAK		
		NDA sessions	**These sessions require invidividuals or their institutions to be covered under proper NDA**	
	10:05	Intel NDA Session	Intel	
	11:10	BREAK		
	11:20	NVIDIA NDA Session	NVIDIA	
	12:25	**LUNCH (on your own)**		
Session chair:		**Apps / optimizations / algorithms**		**Application/algoroithm and/or platform-specific optimizations**
Rebecca Hartman-Baker	13:45	Yeom	LLNL	Data-dependent Performance Modeling of Linear Solvers for Sparse Matrices
	14:00	Ferenbaugh	LANL	Coarse vs. fine-level threading in the PENNANT mini-app
	14:15	Parker	ANL	Performance Optimization and Portability of the Nekbone Mini-App
	14:30	Garrett	LANL	A first look at optimizing performance on the KNL
	14:45	Morozov	ANL	Portability of HACC - a highly tuned cosmology application
	15:00	BREAK		
	15:15	Keipert	ANL	Experiences and challenges while modernizing GAMESS for Theta and Aurora
	15:30	Rennich	NVIDIA	GPU Performance Optimization of the Sweep Operation in Kripke
	15:45	Joo	JLab/ANL/LBL	Experiences and Challenges for Performance Portability in Lattice QCD
	16:00	Vazquez-Mayagoitia	ANL	Many-core and GPU developments in the parallel ELectronic Structure Infrastructure library (ELSI)
	16:10	BREAK		
Session Chair:		**Performance Portable Abstractions**		**General abstractions suitable for managing portability in multiple applications**
Hai Ah Nam	16:40	Nguyen	LBL	Portable Data Locality Management with High-Level Programming Abstractions
	16:55	Vetter	ORNL	Understanding Portability of a High-Level Programming Model on Diverse HPC Architectures
	17:15	Trott	SNL	Kokkos - Performance Portability Today
	17:35	Hornung	LLNL	The RAJA Encapsulation Model for Architecture Portability
	17:55	Jacob	IBM	Towards Performance Portable GPU Programming with RAJA
	18:15	ADJOURN (dinner on your own)		

Day Two

	7:30	Coffee/Light Breakfast		Mingle
	8:15	Opening Remarks		Welcome, recap of day 1, overview of day 2, ….
Session Chair:		**Managing the Memory Hierachy**		**Abstractions/techniques for managing data motion between standard DRAM and HBM/Device memory**
TBD	8:20	Poliakoff	LLNL	Copy Hiding Application Interface (CHAI): Hiding Data Motion for Performance Portability
	8:30	Sakharnykh	NVIDIA	Harnessing Performance of Geometric Multi-Grid Methods by using LOC and TOC architectures
	8:45	Delalondre	ANL	Leveraging heterogeneous systems and deep memory hierarchies for brain tissue modeling
	9:05	DeRose	Cray	Cray's Prog. Env. for Portable Performance and Programmability on Systems with High-Bandwidth Memory
	9:20	Karlin	Multi-lab	Quad Lab Proposal of Fundamental Cross Architecture Multi-Level Memory Support
		Application Experience with Performance Portable Abstractions		
Session Chair:	9:40	Kim	IBM	An abstraction for unstructured mesh problems
Tjerk Straatsma	9:55	Kunen	LLNL	Nested Loop RAJA for Performance Portability
	10:10	Moore	SNL	Obtaining Threading Performance Portability in SPARTA using Kokkos
	10:25	BREAK		
	10:55	Beckingsale	LLNL	Lightweight Models for Dynamically Tuning Data-Dependent Code
	11:05	Womeldorff	LANL	Kokkos and Legion Implementations of the SNAP Proxy Application
	11:15	Bleile	LLNL	Investigation of Portable Event-Based Monte Carlo Transport
	11:30	Grinberg	IBM	Performance portable single source-code implementation of sparse linear algebra operations on CPUs and GPUs
	11:45	Peles	LLNL	Investigating interoperability and performance portability of select LLNL numerical libraries
	12:05	Pennycook	Intel	Performance Portability of Kernel-based Abstractions
	12:25	LUNCH (provided / breakout topics by table)		
Breakout #1	13:20	**BREAKOUT SESSION #1 (Managing the Memory Hierarchy / Performance Portable Abstractions)**		
		Breakout Leads:		Doug Doerfler, TBD (mem hierarchy)
				Brian Friesen, TBD (PP abstractions)
Session Chair:		**OpenMP**		**Experience with OpenMP and recommendations on guiding future standards**
Hai Ah Nam	14:50	Pennycook	Intel	Generalizing a DSL for Structured Dependency (Stencil-like) Codes to OpenMP Loops
	15:10	Levesque	Cray	How we can get Hybrid OpenMP/MPI to out perform All-MPI
	15:30	Bertolli	IBM	Performance Portability with OpenMP on Nvidia GPUs
	15:50	Larkin	NVIDIA	Performance Portability Through Descriptive Parallelism
	16:10	Martineau	UK	Investigating the performance portability capabilities of OpenMP 4, Kokkos and Raja
	16:30	BREAK		
	17:00	Appelhans	IBM	Performance Portability Experience with LLVM, OpenMP 4, and Kripke
	17:15	Eichenberger	IBM	OpenMP Specifications for Portability
	17:30	Hernandez	ORNL	Experiences with High-Level Programming Directives for Porting SPEC ACCEL on multiple architectures
	17:45	Scogland	LLNL	Performance Portability with OpenMP: Experiences with 4.5 and Looking Toward 5.0
	18:05	Adjourn (dinner on your own)		
	19:30 - 22:00	Intel NDA Session		Optional set of evening talks on Intel NDA material for interested attendees

Day Three

	7:30	Coffee/Light Breakfast		Mingle
	8:15	Recap of breakout #1		Each of four groups to present 8-10 minute summary
Session Chair:		**Tools / Compilers**	**Tools for performance portability and analysis**	
Hai Ah Nam	9:00	Cook	SNL	The Importability of Performance Tools
	9:10	Gonzalez	IBM	Next-gen profiling-infrastructure for supercomputers based on hybrid nodes
	9:20	Laguna	LLNL	STATuner: Tuning CUDA Kernels via Compiler Analysis and Machine Learning
	9:35	Hammond	SNL	Profiling Interfaces for Parallel C++ Abstractions - KokkosP
	9:50	Basu	LBL	Leveraging Compiler-Based Tools for Performance-Portability
	10:10	Poxon	Cray	Adding Parallelism to HPC Applications using Reveal
	10:25	BREAK		
Session Chair:		**IO / Burst Buffers**	**The I/O bottleneck and use of burst buffers**	
TBD	10:55	Miller	LLNL	Probing Portable Performance of Parallel I/O Paradigms using MACSio
	11:10	Ovsyannikov	LBL	ChomboCrunch and VisIt for carbon sequestration and in-transit data analysis using burst buffers
	11:30	Mohror	LLNL	Performance Portability for Burst Buffers with the Scalable Checkpoint / Restart Library (SCR)
Session Chair:		**Domain Specific Languages**		**Use of DSL's for performance portability**
TBD	11:50	Richards	LLNL	Portable Performance in Real Applications using Generated Code
	12:05	Straalen	LBL	AMRStencil: An Embedded DSL for Expressing Structured Adaptive Mesh Refinement Algorithms
	12:20	Ibrahim	LBL	Performance Portability Through Unifying the Interface to Multiple Programming Models
	12:35	LUNCH (provided / breakout topics by table)		
Breakout #2	13:30	**BREAKOUT SESSION #2 (OpenMP Futures, Tools/Compiler/System Requirements)**		
		Breakout Leads:	TBD, TBD (OpenMP Futures)	
			TBD, TBD (Tools/Compilers/System)	
	15:00	BREAK		Scribes for breakout given time to collect notes
	15:30	**Recap of breakout #2**		Each of four groups to present 8-10 minute summary
		Wrapup discussions		
	16:15	Vendor Q&A / Panel		Vendor reps to discuss challenges and answer Q&A
	16:55	Wrapup / next-steps / takeaways		Capture followup goals, decide on subsequent meetings and potential topics
	17:15	ADJOURN		
	17:55	DINNER		